MOSIG – MSIAM – 2020-2021 Information Access and Retrieval – GBX9MO23
Georges Quénot – Philippe Mulhem – Jean-Pierre Chevallet
1st February 2020 – 9h00-11h00 (9:00am-11:00am) – 2 hours

Course materials, the two papers related to the examination, personal notes, and calculators (without network capabilities) are allowed. ***No Internet network connection or other communication means should be used during the examination except for downloading the subject or for asking questions about the subject if necessary via the zoom room dedicated to that.***

Your scanned or photographed copies should be sent to Georges.Quenot@imag.fr and to Jean-Pierre.Chevallet@imag.fr as soon as possible after 11am, normally within half an hour. We will understand if you encounter difficulties. Let us know as soon as possible if this is case. **Please check that they are well readable** after scanning (preferred) or photographing.

You can alternatively use a text editor for generating and sending a PDF **but**, in this case, y**ou are not allowed to copy and paste anything from the papers, from the course material or from anything else in the document,** and you should send your copies right after 11am.

Don't forget to put your name and your student number on your exam copies, even if you send them by mail. Also, please send them with your UGA mail account.

The examination consists in questions related to two scientific papers and/or to the contents of the course:

[1] F. Faghri, D. J. Fleet, J. R. Kiros, S. Fidler, "VSE++: Improving Visual-Semantic Embeddings with Hard Negatives", Proceedings of the British Machine Vision Conference (BMVC), 2018. (BMVC Spotlight)

[2] Zamani, H., Dehghani, M., Croft, W. B., Learned-Miller, E., & Kamps, J. (2018, October). "From Neural Re-Ranking to Neural Ranking: Learning a Sparse Representation for Inverted Indexing", Proceedings of the 27th ACM international conference on information and knowledge management (pp. 497-506).

In the following questions, we expect real explanations with details, and not only excerpts from the papers. You should spend about 5 minutes per question and we expect concise answers.

Questions related to paper 1.

Q1.1: What retrieval tasks are considered in this paper?

Q1.2: What is a "Visual-Semantic Embedding"?

Q1.3: What is the main contribution in the work described in this paper?

Q1.4: What is a "hard negative" in the context of this paper? Why is it considered?

Q1.5: a) What is the difference between the "Sum Hinge" and the "Max Hinge" losses? b) Which is better and why? c) Compare their computing runtime complexities.

Q1.5: a) What is the difference between the "VSE" and the "order" losses? b) Which is better and why? c) Compare their computing runtime complexities.

Q1.7: What are inconveniences of choosing too small or too large mini-batch sizes?

Q1.8: a) List the various improvements considered between the simplest and the most complex versions of the considered system. b) Rank them according to the gain that they each brings in terms of reported top-1 accuracy.

Q1.9: Considering the 8 results displayed in Figure A.1, in which proportions does VSE++ performs better, equally well or worse than the VSE0 baseline?

Q1.10: a) What are the main differences between the VGG-19 and the ResNet-152 architectures? b) What do they have in common? c) Which performs better in the considered experiments and why?

Q1.11: What is the role of the α parameter in equations (5) and (6)?

Q1.12: What alternatives to the "Sum Hinge" and the "Max Hinge" losses could be considered?

Questions related to paper 2.

Q2.1: IR systems are supposed to solve a query using natural language or structured queries. Explain why a classical database is usually not suited to implement an IR system, and briefly describe what is used instead.

Q2.2: Explain with your own words what is "relevance feedback" in RI, then justify why most IR systems (including the one presented in this paper) are using pseudo relevance feedback instead of "real" relevance feedback.

Q2.3: Explain with your own words, what a "dense representation" is and why this representation is preferred in Neural Based IR instead of words.

Q2.4: In previous query (2.3), we claim a "dense representation" is preferred an input for a Neural Network (NN), but can an IR system using NN could use a non-dense representation for word? If yes explain how to do it; if no justify why it is impossible.

Q2.5: Explain with your own words why most neural ranking system can only re-rank document and describe what is used for initial ranking.

Q2.6: Explain the notion of Zipfian distribution and justify why a dense representation is not a Zipfian distribution.

Q2.7: Explain with your own words why sparsity of a query is important in this IR system.

Q2.8: In the equation (2) of this paper, justify the role of the "+ 1" in the normalization and in the sum. Explain what exactly this sum is about and propose a complexity formula of this calculation using paper notations (i.e. same symbols).

Q2.9: Authors claim that equation (2) ensure a direct link between the length of document and the computed output. The justification of this claim is not detailed in this paper. First explain how they compute document length, and then justify this author claim, finally explain how to insure that the query representation will be sparser than document one.

Q2.10: With your own words, explain the figure (3), and details, why authors can almost be sure that sparsity can be learn with their framework.